Evaluation of techniques for displaying shape and disparity between paired surfaces

Chris Weigle

University of North Carolina, Chapel Hill

May 6, 2002

Abstract

In this paper we present results of an investigation of three techniques for displaying a pair of surfaces such that their shapes and mutual disparity are readily compared. We require the surfaces to be displayed such that they overlap in the user's view and take disparity to mean a difference in height above some reference plane (sufficiently, one behind the user and perpendicular to the user's view). The goal of the visualization is to enable tasks that involve both shape classification and identification of points of extreme (minimum and maximum) surface disparity. The techniques use shading, color, texture, and shadows as cues to shape or disparity – specifically, the first technique displays only the "near" surface and maps color to disparity, the second technique modulates the "near" surface by an opacity texture allowing the "far" surface to be visible where opacity is low, and the third technique builds on the second by adding shadows cast from the opaque portions of the "near" and onto the "far" surfaces. The results of the study do not indicate that any one of the three tested techniques is significantly better for simultaneously enabling the described tasks.

Introduction

The most fundamental goal of any three-dimensional visualization is to convey spatial relationships inherent in the underlying data. Without the availability of special equipment, this goal is confounded by the necessity to project that three-dimensional space to a two-dimensional image for display. Information about the spatial relationships between objects along the lines of projection is, previously encoded directly in the three-dimensional geometry, is lost in the process of projection. Many visualization techniques attempt to restore this lost information through application of cues which, though they may have no physical relationship to the underlying data, are readily understood by the human visual system. Some of these cues, such as illumination and shadows, are everyday indicators of shape or depth.

When objects to be compared project to the same part of the image (typically with one obscuring the other), these cues may not be sufficient. For example, an oncologist may wish to compare a magnetic resonance image of a patient's tumor with the high-dosage region of a proposed radiation treatment. The two geometries need to be compared to ensure the high-dose region is everywhere outside the tumor but never too far into healthy tissue. Information is required for two surfaces which logically and physically overlap. Since it is likely more informative to place both objects in their correct positions (continuing the example, surrounding tissues may constrain where it is most critical the high-dosage region not stray from the tumor), it will be necessary to find a way to present details of both surfaces simultaneously where they project to the same part of the image. Similar visualization problems exist when comparing model or simulation data with that collected during the modelled experiment, comparing multiple runs of a simulation with differing parameters, or comparing geometry of surfaces scanned by different methods, to name a few.

There are many psychological studies showing the ability to perceive shape from shading (Blake and Bulthoff, 1990; Christou and Koenderink, 1997; Curran and Johnston, 1996; De Haan et al., 1995; Gibson, 1950; Ramachandran, 1988). All such shape perception is, however, influenced by our visual systems' bias toward interpreting a scene as light from overhead (Gibson, 1950) – unless strong cues are given to contradict such illumination. Although there is some ambiguity in the perception of shaded images, these ambiguities can be compensated for by adding texture (Curran and Johnston, 1996), specular highlights (Blake and Bulthoff, 1990), cast shadows (Erens et al., 1993), or object boundaries (Ramachandran, 1988), to name a few. Texture is itself a strong cue to surface shape under the right conditions, namely a strong texture component oriented in the direction of principal curvature of the surface and an image of the textured surface formed by projection (Li and Zaidi, 2000, 2001).

The most common illumination model used in computer graphics, the empirical Phong lighting model, may convey shape cues in a similar manner as natural objects under natural illumination (Johnston and Curran, 1996). The Phong lighting model approximates both diffuse and specular lighting according to Lambert's Law and Snell's Law, respectively (Phong, 1975). Certainly, under the right conditions, Phong illumination conveys a sense of shape and depth.

A common visualization technique for showing how a scalar parameter varies over a surface is the colormap. Research has shown how to construct a perceptual colormap to enhance performance of either a metric estimation or shape classification (Healey, 1996; Ware, 1988). These constructed, perceptual colormaps have also been demonstrated (in their isoluminant form) to work on illuminated, textured surfaces. One technique for displaying two objects where one is inside the other is to display the near surface as semi-transparent with curvature-oriented texture elements (Interrante et al., 1996). Curvature-oriented textures conveys both shape and disparity and has been shown to combine well with other common visualization methods.

This paper describes a preliminary study comparing three techniques for displaying surfaces such that two general tasks may be simultaneously accomplished. The tasks are that of identifying the points of extreme surface disparity and identifying the surfaces' shapes. In this paper, we will take disparity to refer simply to a difference in distance to a reference plane – we will typically assume the plane to be orthogonal to the direction of view (equivalently the direction of projection). The first technique uses a single shaded, opaque surface with a disparity colormap implying the second surface. The second technique consists of two shaded surfaces, the "near" surface using curvature-oriented texturing as in Interrante et al. (1996), where the texture modulates the surface opacity such that the "far" fully opaque surface is partially visible. The third technique adds shadows cast from the textured "near" surface of technique two onto the opaque "far" surface.

Method

Respondents

The participant pool was composed of 25-30 undergraduate psychology students (Psyc 10) from the University of North Carolina, Chapel Hill. These numbers exclude two participants determined to be significant outliers under the color condition, one participant who reported not fully understanding direction until near the end of the run, and one participant with improperly generated samples. The pool of usable participants was composed of 7 females and 12 males with a median age of 19 – females did not participate uniformly across the visualization conditions. The remaining participants were divided six/six/seven between the three visualization conditions (one/four/two were female).

The pool was expected to have only minimal experience with computer graphics used for visualization. The males typically had moderate video game experience; all participants had experience with recent computer-generated motion pictures. Visualization typically utilizes higher geometric detail and lower texture detail than video games and less "realism" than motion pictures – participants' previous experiences were sufficiently different from the computer graphics used in the study.

Participants were self-selected against color-blindness. The data suggest that the two outliers in the color condition were either unaware of their color-blindness or were uncertain of the instructions.

Participants were given a one hour credit towards the completion of the Psyc 10 experiment participation requirement (a total of six hours are required per student).

Materials

All measures were conducted on an Intel Pentium4 1.7GHz PC, a 21" Sony Trinitron CRT, and an NVIDIA GeForce2 graphics card. Measurements took place in a room with controlled lighting conditions. The monitor was driven at an 85Hz refresh rate, at a resolution of 1600x1200, with the participant at a normal desktop computer viewing distance (eighteen to twenty-four inches) such that the image of the surfaces typically subtended about 20 degrees of visual angle. Custom software was used to display the surfaces, present the tasks, and accept and record the participants' responses.

Participants used a pointing device (a mouse) to respond to questions about the location of extremal of disparity by clicking the appropriate location on the "near" surface with the pointing device. Responses to questions about surface shape were indicated by manipulating a slider adorned with representative shapes (the top, middle, and bottom of the slider represented convex, saddle, and concave shapes respectively). Participants were also allowed to view the surfaces under computer controlled animation – the surfaces rocked left and right 30 degrees about the vertical axis – as often as they wanted.

The pairs of surfaces were preselected and randomly ordered. The pairs were chosen to encompass three surface types (convex, concave, and saddle) with relatively low curvatures to constrain the disparity from growing too large and to control the projection of the shadows in the shadow condition. The surfaces were also translated and rotated (in the image plane) to more nearly approximate real situations.

There were two kinds of task: identify the extremal point, and identify the surface shape. Participants would perform each task for each surface pair, either locating (randomly) the minimum or the maximum point of disparity and classifying the shape of either the "far" or the "near" surface. In the color condition, the "far" surface must be constructed from the shape of the displayed "near" surface and the color map. Examples from each condition are included as Appendix A.

Procedure

- Participants read and signed the consent form (Appendix B).
- Participants were given instructions on their visualization condition (Appendix C), the tasks to be performed, and the user interface.
- Participants had a brief (six surface pairs) training session where the program would show the correct responses to the training surfaces.
- Participants performed the tasks for the 32 randomly ordered real test surface pairs.
- Participants were debriefed.
- Participants were given their Psyc 10 credit and receipt.

Results

Surface Disparity

Participant disparity responses were transformed from raw surface disparity to an error value as $error = |true\_value - participant\_response|$ for minimum and maximum cases. Summary statistics for the transformed estimates appear in Figure 1. So converted, the participants' estimated disparity error had significant predictors as noted in the ANOVA table in Figure 2.

min/max vizid	2-Way Tables of Descriptive Statistics N=608 (No missing data in dep. var. list) est. disp. error Means	est. disp. error N	est. disp. error Std.Dev.
0	0.122235	304	0.159387
0	0.065299	96	0.083338
1	0.174961	96	0.176315
2	0.125844	112	0.177694
1	0.216636	304	0.269022
0	0.064933	96	0.144113
1	0.240932	96	0.247693
2	0.325842	112	0.307940
All Groups	0.169436	608	0.225919

Figure 1: Summary statistics for estimated disparity error.

Effect	Univariate Tests of Significance for est. disp. error Sigma-restricted parameterization Effective hypothesis decomposition SS	Degr. of Freedom	MS	F	p
Intercept	16.72655	1	16.72655	395.7294	0.000000
"min/max"	1.18516	1	1.18516	28.0393	0.000000
vizid	3.08692	2	1.54346	36.5163	0.000000
"min/max"*vizid	1.09431	2	0.54716	12.9451	0.000003
Error	25.44512	602	0.04227		

Figure 2: ANOVA table for predictors – visualization condition (vizid: color = 0, texture = 1, shadow = 2); disparity query (min/max: min = 0, max = 1).

Figure 3 shows the results of Tukey tests for least significant difference for this data. Figure 4 shows how the predictor interactions relate to the measured participant error.

			Tukey HSD test; Variable: est. disp. error Marked differences are significant at p < .05000					
min/max	vizid		{1} M=.06530	{2} M=.17496	{3} M=.12584	{4} M=.06493	{5} M=.24093	{6} M=.32584
0	0	{1}		0.003012	0.278161	1.000000	0.000020	0.000020
0	1	{2}	0.003012		0.519945	0.002874	0.226953	0.000022
0	2	{3}	0.278161	0.519945		0.271625	0.000818	0.000020
1	0	{4}	1.000000	0.002874	0.271625		0.000020	0.000020
1	1	{5}	0.000020	0.226953	0.000818	0.000020		0.035371
1	2	{6}	0.000020	0.000022	0.000020	0.000020	0.035371	

Figure 3: Tukey tests for least significant difference across significant predictors of estimated disparity error.

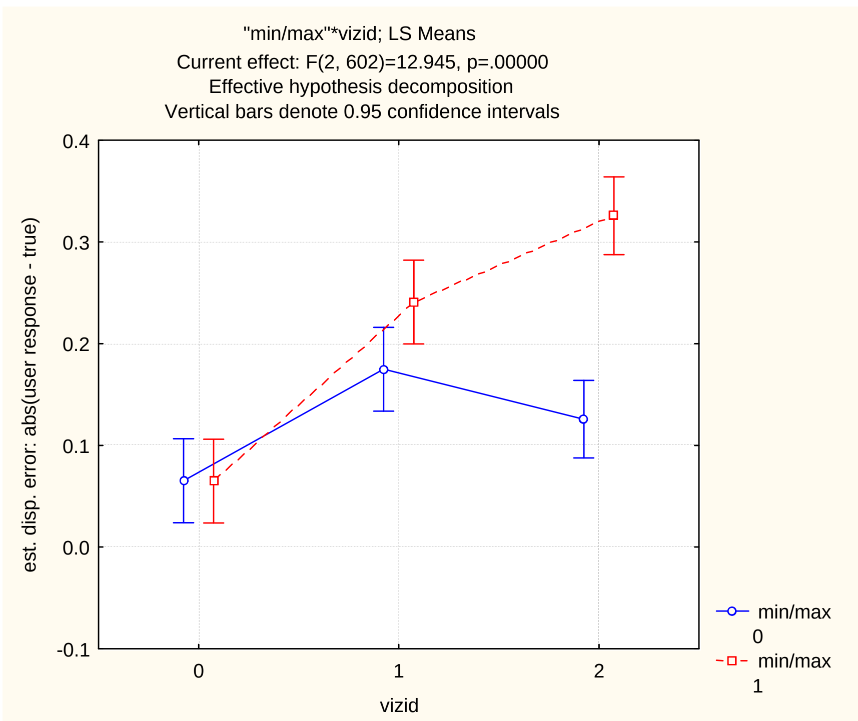

Figure 4: Plot of error versus visualization type for each query type. Error improves for the minimum query but increases significantly for the maximum query.

The data show that visualization condition, and query type predict participant error ($p <$

0.001). The color condition shows the most accurate responses, which was expected. Shadow is in between color and texture for the minimum query (also expected) but significantly worse than texture for maximum. This was not expected, and may be attributable to the scale of the maximum disparity values. The minimum distance between surface pairs is fixed at 0.2 units (a complete surface is a 2.0x2.0 unit patch, ignoring the curved corners) while the maximum curvature ranges from 1.0 to 2.0 units. So the maximum disparities are both an order of magnitude larger than the minimum, but are also on the order of the full extent of the surface patches. Also, the shadow condition may break-down for the maximum disparity case due to the distance the shadow is cast away from the occluding texture.

Surface Shape

The raw data for participant response to the shape query are real numbers indicating position of the slider interface (+1 indicated convex, 0 indicates saddle, -1 indicated concave). The three distributions are not readily compared, unfortunately – since it is not possible or meaningful to respond beyond the range of the slider (-1 to +1), only saddle yields a normal distribution and the other two are severely skewed. The data was therefore transformed into a simple binomial – if the absolute difference between the correct answer and the participant response is less than $\frac{1}{3}$ then the binomial evaluates to one (true) otherwise zero (false). This is a slightly conservative conversion as the error values between $\frac{1}{3}$ and $\frac{1}{2}$ could also be interpreted to indicate a response of the particular shape in question; in practice it did not make a meaningful, significant difference.

The transformed data effectively represents the probability that a participant correctly answered a shape query. In the following figures this probability is referred to as $ESE < \frac{1}{3}$ – the probability that the estimated shape error is less than $\frac{1}{3}$. Figures 5 and 6 gives the summary statistics for the near and far target queries, respectively, and Figure 7 gives the

ANOVA table for the expected predictors.

vizid true shape	2-Way Tables of Descriptive Stats Within: near/far:0 N=608		
	ESE < 1/3 Means	ESE < 1/3 N	ESE < 1/3 Std.Dev.
0	0.549020	102	0.500049
-1	0.407407	54	0.495966
0	0.666667	30	0.479463
1	0.777778	18	0.427793
1	0.705882	102	0.457895
-1	0.592593	54	0.495966
0	0.833333	30	0.379049
1	0.833333	18	0.383482
2	0.361345	119	0.482421
-1	0.206349	63	0.407935
0	0.571429	35	0.502096
1	0.476190	21	0.511766
All Groups	0.529412	323	0.499909

Figure 5: Summary statistics for ESE (estimated shape error) probability for the near surface.

vizid true shape	2-Way Tables of Descriptive Stats Within: near/far:1 N=608		
	ESE < 1/3 Means	ESE < 1/3 N	ESE < 1/3 Std.Dev.
0	0.466667	90	0.501683
-1	0.333333	48	0.476393
0	0.666667	18	0.485071
1	0.583333	24	0.503610
1	0.555556	90	0.499688
-1	0.416667	48	0.498224
0	0.722222	18	0.460889
1	0.708333	24	0.464306
2	0.619048	105	0.487950
-1	0.464286	56	0.503236
0	0.809524	21	0.402374
1	0.785714	28	0.417855
All Groups	0.550877	285	0.498280

Figure 6: Summary statistics for ESE probability for the far surface.

Figure 8 shows that the probability of receiving a correct shape response from the partic-

Effect	Univariate Tests of Significance for ESE < 1/3 Sigma-restricted parameterization Effective hypothesis decomposition				
	SS	Degr. of Freedom	MS	F	p
Intercept	182.9217	1	182.9217	834.9192	0.000000
"near/far"	0.0243	1	0.0243	0.1107	0.739454
vizid	1.6829	2	0.8415	3.8407	0.022017
true shape	13.2247	2	6.6124	30.1812	0.000000
"near/far"*vizid	4.3365	2	2.1683	9.8967	0.000059
"near/far"*true shape	0.0470	2	0.0235	0.1073	0.898316
vizid*true shape	0.1490	4	0.0372	0.1700	0.953677
"near/far"*vizid*true shape	0.2444	4	0.0611	0.2789	0.891634
Error	129.2626	590	0.2191		

Figure 7: ANOVA table for predictors – visualization condition (vizid: color = 0, texture = 1, shadow = 2); shape query (true shape: convex = 1, saddle = 0, concave = -1); shape query target surface (near/far: near = 0, far = 1).

ipants was significantly lower for concave surfaces. Figure 9 shows that the shadow condition performed unexpectedly poorly at conveying the shape of the near surface.

Visualization condition and the true shape are shown to be significant predictors of the probability of correctly identifying the surface shapes ($p < 0.05$). There also exists a significant interaction term between visualization type and whether the shape target is the near or far surface ($p < 0.001$). This interaction term predicts the relatively poor performance of the shadow condition at conveying the shape of the "near" surface. Because of the similarity between the texture and shadow conditions in displaying the "near" surface, this suggests that the shadow condition is too sparse to allow the participant to fill-in the overall shape of the surface.

The shadow condition also performs about ten percent better at conveying the "far" surface. This could be due to additional shape cues from the shadows, but following the arguments for the "near" surface problems, may more likely be due to obstructing less of the view of the "far" surface.

The data also show that the concave surfaces caused significant difficulty for identification,

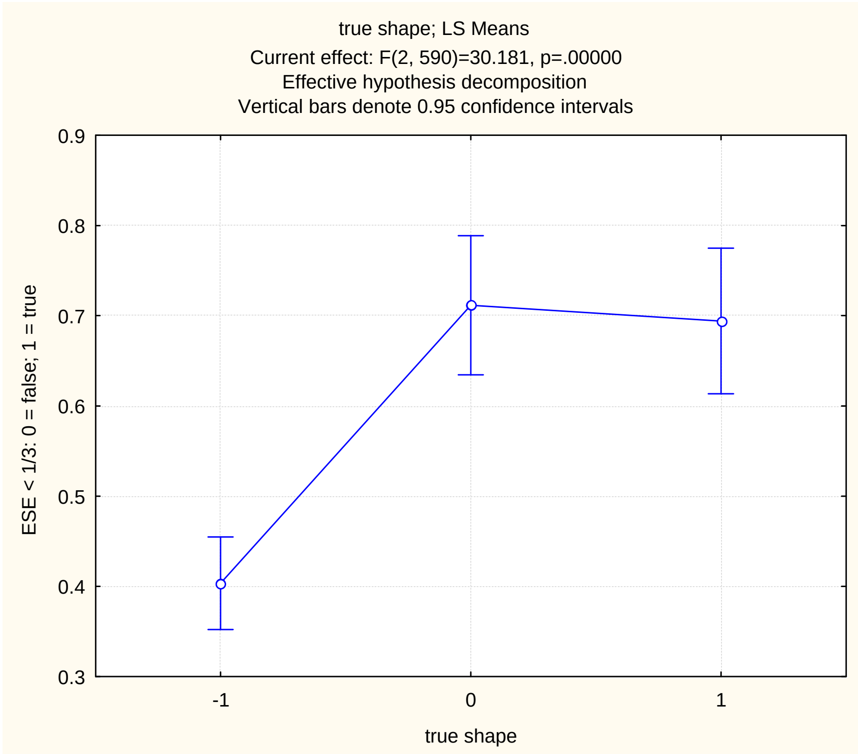

Figure 8: Plot of probability versus true shape. Participants were unable to recognize concave surfaces.

regardless of visualization condition or which surface was the target. This is most likely due to the illumination of the surface's – it may be too close to the horizon and not sufficiently overhead to take advantage of our internal representation of illumination direction.

Discussion

The shadow visualization condition performed as hoped when identifying points of minimum disparity, but did not perform well at maximum disparity. This is likely due to the design of the surface geometry, and deserves to be retested before the poor performance at identifying maximum disparity is taken as a property of the technique.

The shape queries expose a design flaw in the illumination of the surfaces. The apparent

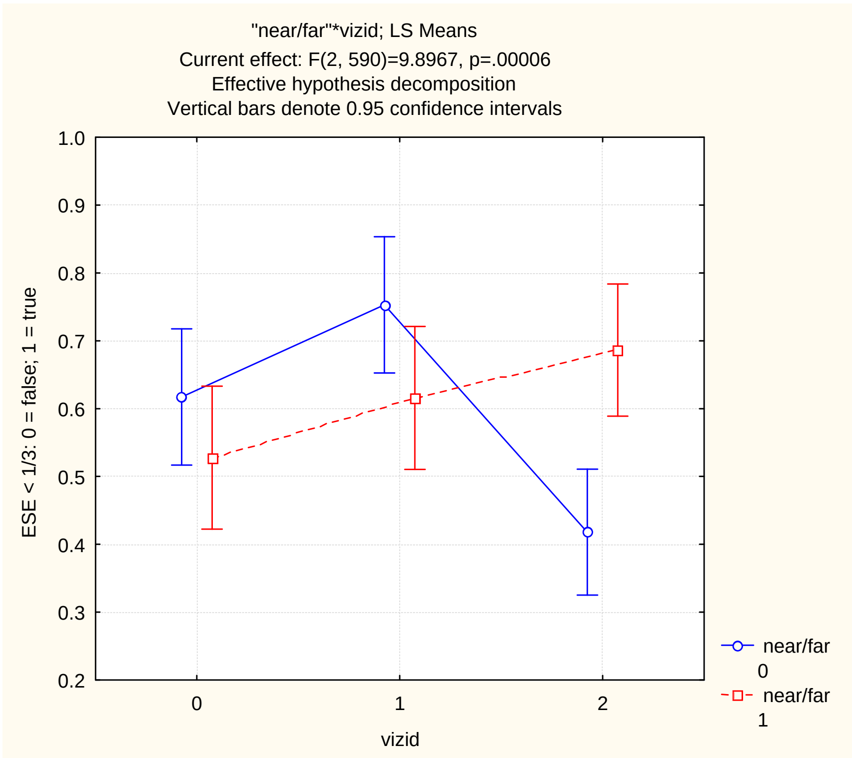

Figure 9: Plot of probability versus visualization condition. Note that the shadow condition shows poor performance for recognizing the near surface.

inability of the participants to distinguish the concave shapes indicates that the light source was positioned too far from the human visual system's internal "overhead" position – there are not enough other cues in the visualization techniques to overpower that internal model and the concave surfaces were not perceived as such accordingly. The shadow condition did, otherwise, enable some improved recognition of "far" shapes over the other two conditions.

Immediate future work should be to redesign the geometry, reposition the light source, and retest. Also suggested by the data is the need to explore the texture density to strike a balance between coverage of the "near" surface and visibility of the "far" surface (and of the shadows in the shadow condition).

References

Blake, A. and Bulthoff, H. (1990). Does the brain know the physics of specular reflection? *Nature*, 343:165–168.

Christou, C. G. and Koenderink, J. J. (1997). Light source dependence in shape from shading. *Vision research*, 37(11):1441–1449.

Curran, W. and Johnston, A. (1996). The effect of illuminant position on perceived curvature. *Vision research*, 36(10):1399–1410.

De Haan, E., Erens, R. G. F., and Noest, A. J. (1995). Shape from shaded random surfaces. *Vision research*, 35(21):2985–3001.

Erens, R., Kappers, A., and Koenderink, J. (1993). Perception of local shape from shading. *Perception and Psychphysics*, 54:145–147.

Gibson, J. J. (1950). *The perception of the visual world*. Houghton Mifflin, Boston.

Healey, C. (1996). Choosing effective colours for data visualization. In *Proceedings Visualization '96*, pages 263–270. IEEE Computer Society.

Interrante, V., Fuchs, H., and Pizer, S. (1996). Illustrating transparent surfaces with curvature-directed strokes. In *Proceedings Visualization '96*, pages 211–218. IEEE Computer Society.

Johnston, A. and Curran, W. (1996). Investigating shape-from-shading illusions using solid objects. *Vision research*, 36(18):2827–2835.

Li, A. and Zaidi, Q. (2000). Perception of three-dimensional shape from texture is based on patterns of oriented energy. *Vision Research*, 40:217–242.

Li, A. and Zaidi, Q. (2001). Information limitations in perception of shape from texture. *Vision Research*, 41:1519–1534.

Phong, B. T. (1975). Illumination for computer generated pictures. *Communications of the ACM*, 18:311–317.

Ramachandran, V. (1988). Perceiving shape from shading. *Scientific American*, 259:58–65.

Ware, C. (1988). Color sequences for univariate maps: Theory, experiments and principles. *IEEE Computer Graphics and Applications*, pages 41–49.

Appendix A – Condition Examples

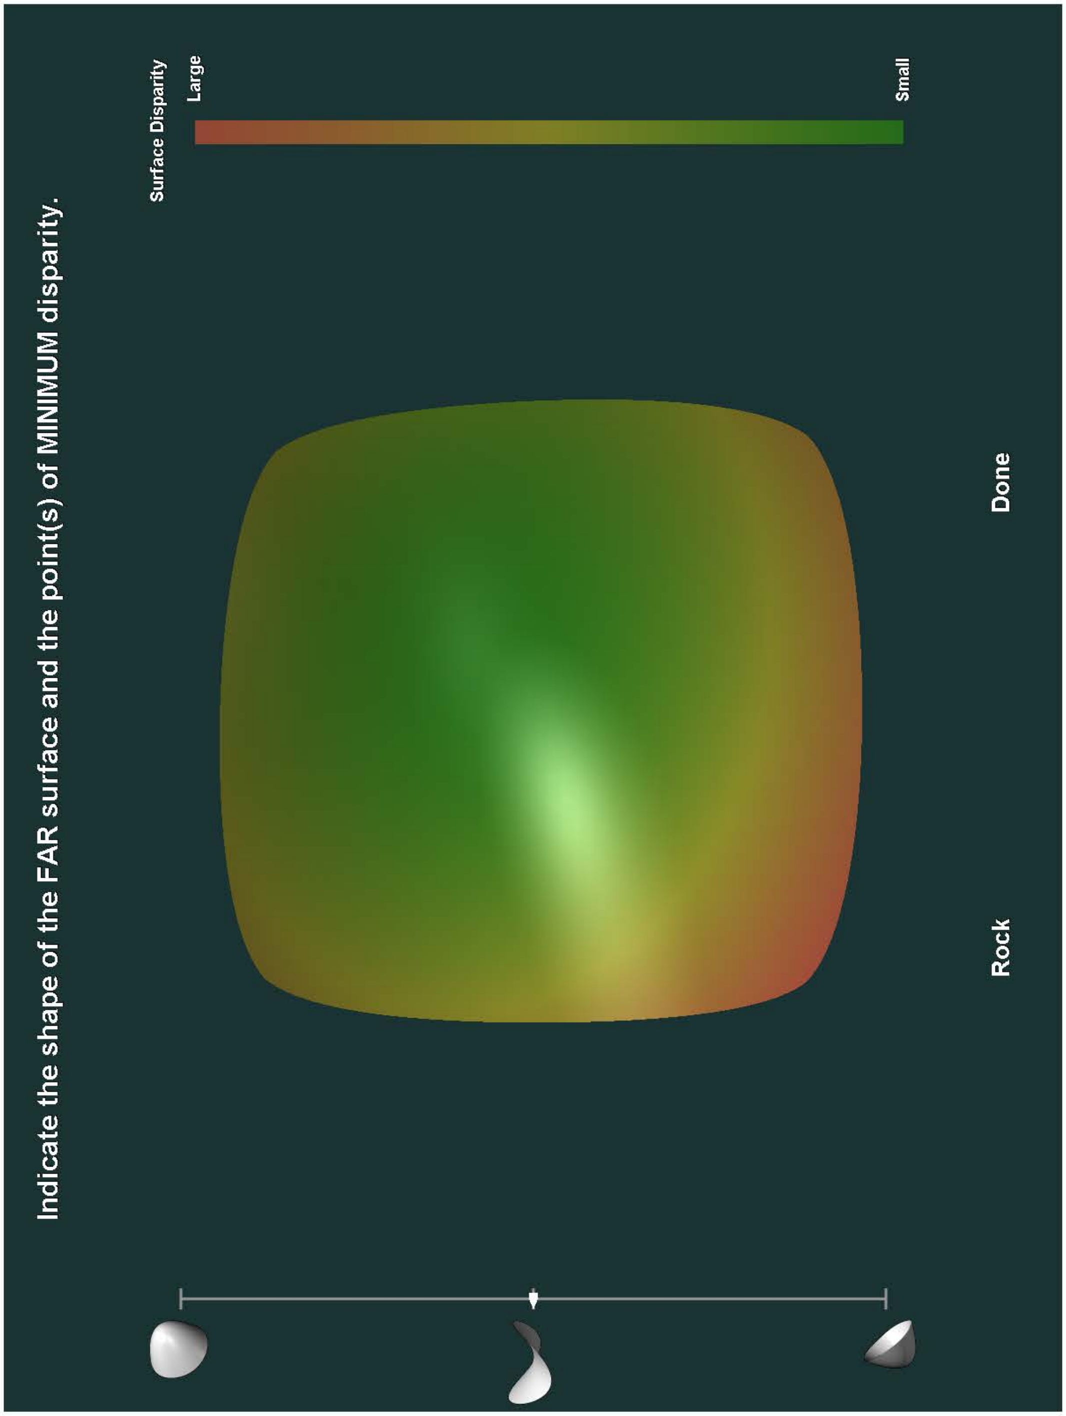

Figure 10: Color Condition

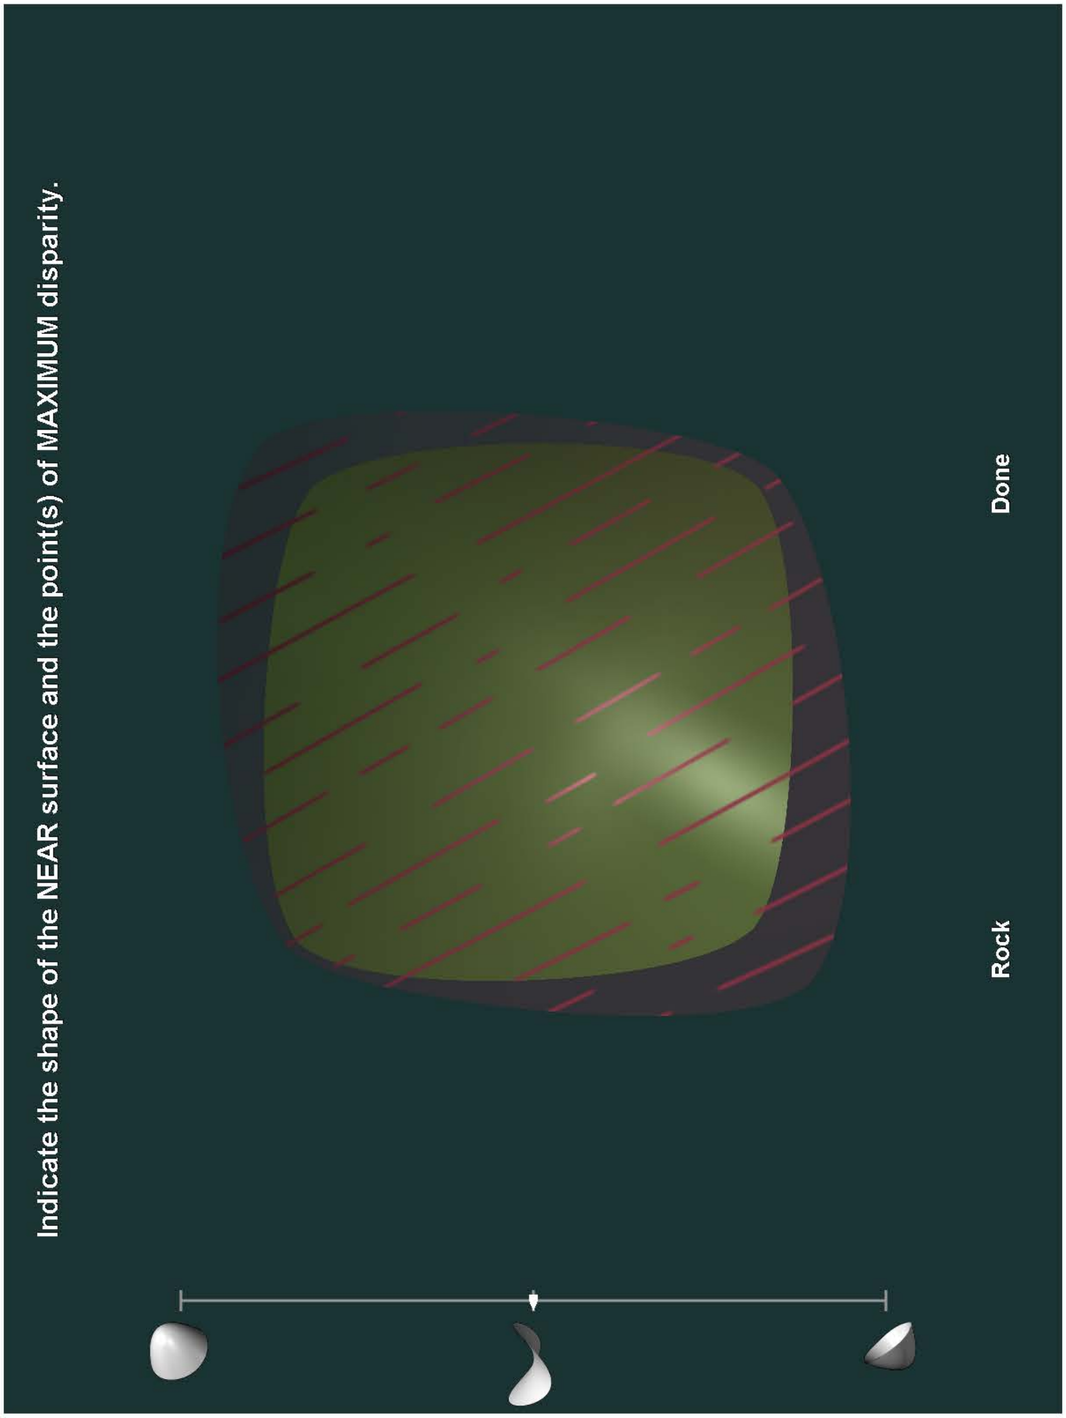

Figure 11: Texture Condition

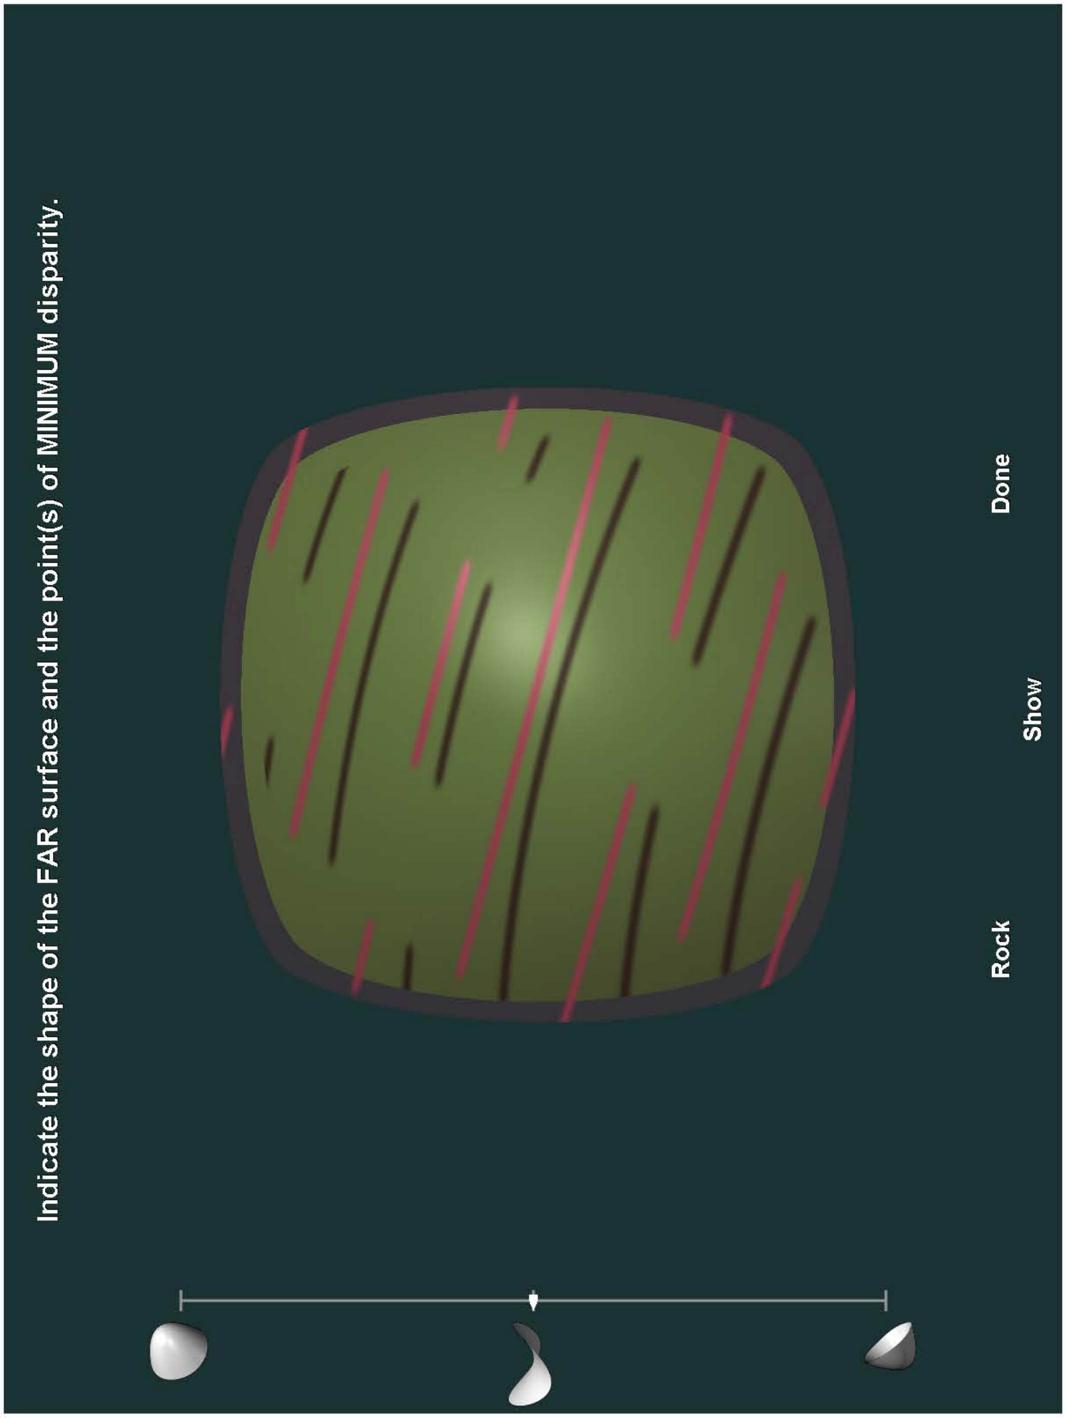

Figure 12: Shadow Condition

Appendix B – Consent Form

THE UNIVERSITY OF NORTH CAROLINA
AT
CHAPEL HILL

Department of Psychology
College of Arts & Sciences

The University of North Carolina at Chapel Hill
CB# 3270, Davie Hall
Chapel Hill, NC 27599-3270

Computer Graphics for Display of Surface Pairs

Introduction and Purpose:

We are inviting 20 to 40 people to take part in a research study investigating techniques for displaying overlapping pairs of surfaces on a computer to facilitate easy comparison between surfaces. We hope to apply what we learn from this study to create tools for scientists, engineers, and medical professionals allowing them to comprehend displays containing overlapping objects. The study is being conducted by Chris Weigle and supervised by Dr. Abigail Panter, with the departments of Computer Science and Psychology, UNC-CH, respectively.

What You Should Expect:

You will be introduced to one of the display techniques and briefly trained in how to perform tasks in the experiment system. You will then be asked to perform two tasks on each of a series of images created using the display technique. The tasks are all performed with a computer mouse. It typically takes an hour to complete the full series of tasks. At the completion of the series, you will be given the opportunity to ask questions about the experiment and the research. If you have questions or concerns after completion of the study, you should contact Chris Weigle at (919) 962-1865 (email: weigle@cs.unc.edu) or Dr. Abigail Panter at (919) 962-4012 (email: panter@unc.edu).

Compensation:

This study qualifies for one hour toward the research requirement for Psychology 10. Partial credit will be assigned if necessary (should you decide to begin but not complete the session) and will be commensurate with portion of the study completed.

Privacy:

Every effort will be made to protect your privacy. Your name or other personal identifier will not be used in conjunction with the information collected in this study or in any reporting of the study. A list of participants will be kept, but will be used only for assigning credit for the research requirement of Psychology 10. There will be no information kept which could be used to associate your identity with your responses to the experiment tasks.

Risks and Discomforts:

To the best of our knowledge, you should be exposed to no more risk or discomfort than from typical computer use.

THE UNIVERSITY OF NORTH CAROLINA
AT
CHAPEL HILL

Department of Psychology
College of Arts & Sciences

The University of North Carolina at Chapel Hill
CB# 3270, Davie Hall
Chapel Hill, NC 27599-3270

Your Rights:

You take part in this study at your own discretion. Should you choose to begin the study, you have the right to terminate your participation in the study prior to completion of the session. There is no penalty for either deciding not to participate or deciding to withdraw.

Institutional Review Board Approval:

The Academic Affairs Institutional Review Board (AA-IRB) of the University of North Carolina at Chapel Hill has approved this study. If you have any concerns about your rights in this study you may contact the Chair of the AA-IRB, Barbara Davis Goldman, Ph.D., at (919) 962-7761 (email: aa-irb@unc.edu).

Informed Consent:

I understand the above, and have had the opportunity to ask questions, and have had them answered to my satisfaction. I have read the information in this consent form and agree to participate in the study. I understand that one copy of this signed consent form will be kept for the experimenters' records.

________ I am 18 or older.

________ I am under 18, but have a parental consent form on file in the Participant Pool Office (Davie 311).

____ _____/_____/_____
(Signature of Participant) (Date)

Appendix C – Instructions

Instructions for condition 1 (single surface with colormap)

In this experiment, you will be shown a series of images of a single surface, the “near” surface. There is also a second surface, the “far” surface, which you are not directly shown. The “near” surface is colored to indicate distance to the “far” surface. It is like coloring an elevation map, but instead of all elevations having the same reference (sea level) each elevation is in reference to a corresponding point on the “far” surface. A legend is supplied to the right of the surface to help you understand the color scale.

For each image, you will be asked to perform two tasks. The two tasks will be presented together - you may perform them in either order you wish. The first task will be to either indicate the point of minimum (or sometimes the maximum) difference between the “near” and “far” surfaces. Indicate your response by pointing and clicking on the surface with the mouse. The second task will be to indicate what you think is the shape of the “near” (or sometimes the “far”) surface. You indicate your response by manipulating a slider (also with the mouse) to indicate if the surface appears to be convex, like at the top of the slider, concave, like at the bottom, or saddle, like in the middle. At any time, you can click on the “rock” button, and the surface will do a little rocking motion to let you get a better feel for what it looks like. When you’ve answered both parts of the question, click “done”.

First you will be presented with 6 images for practice. Your responses will not be recorded. Instead you can have the computer show you the correct answers and you can learn how the interface works and what the shapes look like.

After the training session is done, the there will be 32 real images for you to performs the tasks on. When you are done, I’ll let you ask whatever questions you have and give you your Psyc 10 credit slip.

Instructions for conditions 2 and 3 (2 surfaces, textured, with and without shadows)

In this experiment, you will be shown a series of images of a pair of surfaces, a "near" surface and a "far" surface. The "near" surface is displayed by the reddish stripes that you can see around in place. Behind the stripes is the green "far" surface.

For each image, you will be asked to perform two tasks. The two tasks will be presented together - you may perform them in either order you wish. The first task will be to either indicate the point of minimum (or sometimes the maximum) difference between the "near" and "far" surfaces. Indicate your response by pointing and clicking on the surface with the mouse. The second task will be to indicate what you think is the shape of the "near" (or sometimes the "far") surface. You indicate your response by manipulating a slider (also with the mouse) to indicate if the surface appears to be convex, like at the top of the slider, concave, like at the bottom, or saddle, like in the middle. At any time, you can click on the "rock" button, and the surface will do a little rocking motion to let you get a better feel for what it looks like. When you've answered both parts of the question, click "done".

First you will be presented with 6 images for practice. Your responses will not be recorded. Instead you can have the computer show you the correct answers and you can learn how the interface works and what the shapes look like.

After the training session is done, the there will be 32 real images for you to performs the tasks on. When you are done, I'll let you ask whatever questions you have and give you your Psyc 10 credit slip.